Ultrascalability combined with extreme availability features and industry-leading performance

Product Guide
October 2010

IBM System x3850 X5 / x3950 X5

Architecture Overview

CONTENTS

Architecture Overview	1
Key Highlights	2
Key Features	3
Key Options	12
x3850/x3950 X5 Images	13
x3850/x3950 X5 Specs	14
The Bottom Line	17
Server Comparison	18
For More Information	20
Legal Information	20

Outstanding performance, superior mainframe-like reliability, and fault-tolerant memory characteristics

IBM® has been designing and implementing servers under the X-Architecture® name since 2001. eX5 technology represents the fifth generation of enterprise servers based on the same design principle IBM began with in 1997: to offer systems that are expandable, offer "big iron" reliability, availability, and serviceability (RAS) features, with extremely competitive price/performance on an Intel® Xeon® processor-based system.

The eX5 technology is primarily designed around three major workloads: database servers, server consolidation using virtualization services, and Enterprise Resource Planning (application and database) servers.

If you're using industry-standard Xeon processor-based servers for running business critical applications, the systems that run these applications need the type of technology designed into IBM's eX5 technology systems. eX5 technology represents a $100M+ investment in designing a flagship offering that can harness the power of 4-socket-and-up 64-bit x86 (x64) Xeon processors. The eX5 family includes two **4U 4-socket** rack-optimized chassis (**x3850 X5** and **x3950 X5**), which are offered in two different scaled configurations up to **8 sockets**, with the architectural capability to scale further. In addition, for those workloads that require the maximum available memory, the x3850 X5 and **x3950 X5** support a 1U **MAX5** memory expansion unit, which adds an additional **32 DIMM slots** *per chassis* (up to **96** DIMMs / **1.5TB** total). If you also need more processor sockets, you can add a *second* x3850 X5 or x3950 X5 to form one **8U 8-socket** server with double the memory (up to **2TB**), HDDs/SSDs, adapter slots, and all the other system resources. IBM **FlexNode partitioning** allows you to logically divide an **8-socket** server into **two 4-socket** servers remotely and automatedly, and then revert back to a single 8-socket server, all without requiring a physical visit to the server.

Maybe you'd like to start out with a 4-socket server and possibly add sockets later, if your needs change. Or perhaps you need *more* than a 4-socket server now, but don't want to get locked into a monolithic 8-socket or larger server—again, in case your requirements change. With IBM, you can stop at 4 sockets, or expand memory or CPU sockets further if needed. And if you require a variety of 4- and 8-socket servers in your data center, you only have to qualify one server for all these workloads. This can save much time and effort and speed deployment. You can also save money on software licensing by virtualizing a 4- or 8-processor server into many VMs, rather than using multiple 2- processor servers. Huge amounts of memory also enable more or larger VMs, and larger databases.

IBM X-Architecture pioneered XpandOnDemand™ ("pay as you grow") scalability, which allows chassis to be simply cabled together to form larger scale-up systems. This capability allows IBM to sell a large SMP (symmetric multiprocessing) system at entry price points. With XpandOnDemand, you can start small and later expand as your needs change, without requiring you to buy more than you need up-front or throw away parts later as you expand.

IBM's eX5 technology-based systems are the ideal solution for scale-up database-serving applications on Microsoft® Windows® with Microsoft SQL Server® or IBM DB2®, as well as Linux® with Oracle or DB2. Database hosting demands ultimate server reliability features, and once installed, they tend to grow and grow, requiring ever greater levels of availability. eX5 servers provide exactly that degree to availability.

Another strong application area for the eX5-based systems is enterprise server consolidation activities workloads, including SAP and Oracle. eX5 systems can offer considerable savings over UNIX® deployments, using our certified solution stacks on either Windows or Linux.. Larger servers need more processor, memory and I/O resources, which make maximum use of any applicable virtual machine software licensing fees and deliver superior system utilization levels. The name of the game in consolidation activities is to deploy the fewest new servers possible and help IT staff manage more images with the same or fewer overall people.

These eX5 servers are designed to protect your data with high performance, high reliability, and high availability. They support the latest **eight**-, **six-**, and **four-core** Intel® **Xeon**™ processors, up to **1066MHz** memory access, and up to **24MB** of **L3** cache to help provide you with the computing power you need to match your business needs and growth. In addition, the x3850 X5 and x3950 X5 use industry-standard **DDR3** memory with **Chipkill**™ **ECC** (Error Checking and

Please see the Legal Information section for important notices and information.

Correcting) protection—for high performance and reliability. For even higher levels of availability, the eX5 servers also offer **Memory ProteXion**™, **memory scrubbing**, **memory rank sparing**, and selectable **memory mirroring**. A **dual-port** integrated high-speed **Gigabit Ethernet** controller with **TOE** (TCP Offload Engine) acceleration is standard. In addition, a dual-port 10 Gigabit Ethernet NIC is standard in most models. The 10GbE solution provides the benefits and flexibility of I/O convergence in a single end-to-end solution. The x3850 offers **seven** high-performance **PCIe** adapter slots (six available, if the 10GbE NIC is installed).

The x3850 X5 and x3950 X5 offer industry-leading scalability, including **4-processor** support (upgradeable to **8 processors** / **64 cores**), up to **1TB** of memory *per chassis* (upgradeable to **2TB** overall), and a choice of up to **eight 2.5-inch** internal **Serial-Attach SCSI** (SAS) or **Serial ATA** (SATA) **hot-swap** hard disk drives with an internal storage capacity of **4TB** *per chassis* (up to **16** drives and **6TB** overall, in a two-chassis configuration), or **sixteen 1.8-inch solid-state drives** (**SSDs**) with a capacity of **800GB** per chassis (upgradeable to **1.6TB**).

All models include a **ServeRAID-BR10i V2** controller, which supports hardware **RAID-0**/**1**/**1E**. This controller is installed in a dedicated RAID slot, and can be upgraded with optional IBM **ServeRAID** SAS RAID controllers, which add up to **512MB** of battery-backed cache, **RAID-10**/**5**/**50** support, and optionally **RAID-6**/**60** and full disk encryption, via **SED** (Self-Encrypting Disk). The **4U** size of the chassis helps you maximize your rack investments. Up to **10** of these chassis can be installed in a single 42U rack, for a total of up to **40** processors, **70** PCIe slots, and **80** HDDs or **160** SSDs, offering an ideal balance of performance, storage and I/O slots per rack. Optional Advanced Connectivity Technology (**ACT)** interconnect cabling helps reduce cable clutter and cost and minimizes installation time when interconnecting many rack-mounted servers.

Standard in both the x3850 X5 and the x3950 X5 is the **Integrated Management Module (IMM)** that enables the user to manage and control the server easily—both locally and remotely. The IMM offers a high level of manageability that is designed to keep costs down and the system up—even when network usage increases. The drop-down **light path diagnostics** panel enables quick servicing of the system if a problem develops. These advanced features help maximize network availability by increasing uptime, as do IBM OnForever™ features, such as **hot-swap**/**redundant HDDs**, **power** and **fans**; **Active Memory**™; **temperature-controlled fans** with **Calibrated Vectored Cooling**™; **IPMI 2.0** support, including **highly secure remote power control** and **Serial over LAN**; as well as **text- and graphics-console redirect over LAN**.

Another improvement with the new generation of X-Architecture is the replacement of legacy BIOS with a new generation **Unified Extensible Firmware Interface** (UEFI). UEFI provides a more intuitive user interface and understandable event logs and better management.

With the inclusion of unique IBM service and support features such as **IBM Systems Director**, **IBM Systems Director Active Energy Manager™** (formerly known as PowerExecutive), and **IBM ServerGuide**™, the x3850/x3950 X5 is designed for maximum uptime.

If you need a balance of high-performance four-socket processing, scalability up to eight sockets, enormous memory expansion, and large I/O capacity in a rack-dense environment, these are the ideal systems.

Key Highlights

eX5 Design Features

eX5 design offers numerous features per chassis to boost performance and reduce product and operating costs:

- Support for up to **four** Intel Xeon processors with **12MB** to **24MB** of **L3** cache (processor dependant), Intel VT and Turbo Boost technology, and Intel EM64T technology.
- Supports **Xeon 8-core**, **6-core** or **4-core** processors, scalable to 8 processors in a two-chassis configuration.
- Support for a **MAX5** memory expansion unit, increasing memory capacity by **50%** for memory-intensive workloads.
- IBM **FlexNode partitioning** allows a physical two-node configuration to be remotely reconfigured by software into logical single-node servers, as needs change or for daily peak and off-peak workloads. For example, FlexNode partitioning provides the flexibility of running the server as *one 8-socket* server at night for heavy-duty batch processing, and as *two 4-socket* servers during the day for separate interactive workloads—all remotely and automatedly, without need for an onsite visit.
- Fast **PC3-10600 DDR3** ECC memory with **Chipkill** error correction, **Memory ProteXion**™ (redundant bit steering), **memory scrubbing**, memory **rank sparing**, and selectable **memory mirroring** protection provides speed and high availability.
- **Seven** high-speed **PCIe Gen 2** adapter slots (six available with 10GbE NIC installed) offer investment protection by supporting high-performance adapters, such as 10Gb Ethernet, Fibre Channel, Converged Network Adapter (CNA), and InfiniBand cards; the **two Active**™ **PCIe hot-plug** slots help increase system availability.

Please see the Legal Information section for important notices and information.

- Up to **4TB** of internal high-speed SAS storage using **eight 2.5-inch** HDDs per chassis.
- **Integrated Management Module** (IMM) standard.
- Optional **ServeRAID** controllers that support **RAID-0**/**1**/**10**/**5**/**50**/**6**/**60** and high-performance SSD drives.

Key Features

Multicore Xeon Processors

The x3850 X5 and x3950 X5 ship with **two** high-performance Intel **Xeon 7500 series** processors, and support up to **four** per chassis (node), allowing you to upgrade as business needs require. The eX5 servers also go a step further by allowing you to add a second node to increase the number of sockets from four to **eight** (using **two** chassis).

The system offers a choice of processor clock rates, cache sizes and power draw:

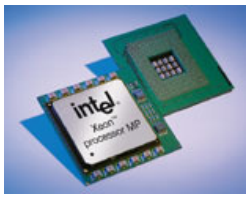

- **130W 8-core** Xeon processor model **X7560** at 2.26GHz, with 64-bit extensions, *low power draw per core* (**16.25W**), **6.4** GTps (gigatransfers per second) QPI speed, **1066MHz** memory access, dual integrated memory controllers, **24MB** of shared **L3** cache, and Intel Turbo Boost Technology and Hyper-Threading technology
- **130W 8-core** Xeon processor model **X7550** at 2.0GHz, with 64-bit extensions, *low power draw per core* (**16.25W**), **6.4** GTps QPI speed, **1066MHz** memory access, dual integrated memory controllers, and **18MB** of shared **L3** cache, and Intel Turbo Boost Technology and Hyper-Threading technology
- **95W 8-core** Xeon processor model **L7555** at 1.86GHz, with 64-bit extensions, *extremely low power draw per core* (**11.88W**), **5.86** GTps QPI speed, **977MHz** memory access, dual integrated memory controllers, **24MB** of shared **L3** cache, and Intel Turbo Boost and Hyper-Threading technology
- **95W 6-core** Xeon processor model **L7545** at 1.86GHz, with 64-bit extensions, *low power draw per core* (**15.83W**), **5.86** GTps QPI speed, **977MHz** memory access, dual integrated memory controllers, **18MB** of shared **L3** cache, and Intel Turbo Boost and Hyper-Threading technology
- **130W 6-core** Xeon processor model **L7542** at 2.66GHz, with 64-bit extensions, *reduced power draw per core* (**21.67W**), **5.86** GTps QPI speed, **977MHz** memory access, dual integrated memory controllers, **18MB** of shared **L3** cache, and Intel Turbo Boost technology
- **105W 6-core** Xeon processor model **E7540** at 2.0GHz, with 64-bit extensions, *low power draw per core* (**17.5W**), **6.4** GTps QPI speed, **1066MHz** memory access, dual integrated memory controllers, and **18MB** of shared **L3** cache, and Intel Turbo Boost Technology and Hyper-Threading technology
- **105W 6-core** Xeon processor model **E7530*** at 1.86GHz, with 64-bit extensions, *low power draw per core* (**17.5W**), **5.86** GTps QPI speed, **977MHz** memory access, dual integrated memory controllers, and **12MB** of shared **L3** cache, and Intel Turbo Boost Technology and Hyper-Threading technology
- **95W 4-core** Xeon processor model **E7520*** at 1.86GHz, with 64-bit extensions, *low power draw per core* (**23.75W**), **4.8** GTps QPI speed, **800MHz** memory access, dual integrated memory controllers, and **18MB** of shared **L3** cache and Hyper-Threading technology

***** This processor is limited to single-chassis configurations (4 sockets per chassis supported).

With the Xeon 7500 series processors, Intel has diverged from its traditional Symmetric Multiprocessing (SMP) architecture to a Non-Uniform Memory Access (NUMA) architecture. The Xeon 7500 processors are connected through three serial coherency links called QuickPath Interconnect (QPI). QPI is capable of **6.4GTps** (gigatransfers per second). Put another way, the 7500 series processors offer **3X** the bandwidth of the previous 7400 series processors.

Turbo Boost Technology dynamically turns off unused processor cores and increases the clock speed of the cores in use, by up to two model frequencies. For example, with six cores active, a 2.26GHz X7560 8-core processor can run the cores at **2.53GHz**. With only three or four cores active, the same processor can run those cores at **2.67GHz**. When the cores are needed again, they are dynamically turned back on and the processor frequency is adjusted accordingly.

Each processor includes two integrated memory controllers, to reduce memory bottlenecks and improve performance. Memory access is at up to **1066MHz** frequency, depending on the processor model and memory used.

In processors implementing **Hyper-Threading Technology**, each core has two threads capable of running an independent process. Thus, an 8-core processor can run **16** threads concurrently.

Intel's **Virtualization Technology** (VT) integrates hardware-level virtualization hooks that allow operating system vendors to better utilize the hardware for virtualization workloads.

Intelligent Power Capability powers individual processor elements on and off as needed, to reduce power draw.

DDR3 Registered Memory with Active Memory Protection

eX5 servers use registered double data rate III (DDR3) DIMMs and provide Active Memory features, including advanced **Chipkill** memory protection, for **up to 16X** better error correction than standard ECC memory, IBM **Memory ProteXion**, and optional **memory mirroring**. In addition to offering **triple the memory bandwidth** of DDR2 or fully-buffered memory, DDR3 memory also uses less energy. DDR2 memory already offered up to **37%** lower energy use than fully buffered memory. Now, a generation later, DDR3 memory is even more efficient, using **10-15% less energy** than DDR2 memory.

The x3850 X5 and x3950 X5 servers support up to **1TB** of memory per node (chassis) in **64** DIMM sockets. They use **PC3-10600** double data rate 3 (DDR3) memory (operating at **1066MHz**) for fast access. Adding a *second* x3850/x3950 X5 server to the first doubles the memory capacity to **2GB**/**128** DIMM slots in **8U** of rack space.

The standard configuration includes **2** memory cards, which support up to **8** DIMMs apiece. The system is upgradeable to **8** memory cards per chassis. (With configurations of 8-to-16 DIMMs, using two memory cards saves cost, but using eight cards increases performance.) If 1TB is not enough, but you don't need additional processors, you can attach a 1U **MAX5** memory expansion unit, which adds another **32 DIMM slots** (up to **512GB**), for a total of **1.5TB** in only **5U**.

Redesign in the architecture of the x7500 series processors brings radical changes in the way memory works in these servers. For example, the Xeon 7500 series processors **integrate two memory controllers inside each processor**, resulting in **eight** memory controllers in a four-socket system. Each processor has four memory channels.

This advanced memory architecture provides up to **333%** more aggregate memory bandwidth (up to **120GBps** when using eight memory cards vs. a maximum of **32GBps** bandwidth) than the previous generation, for exceptional memory performance, and **quadruple** the system memory capacity of the predecessor x3850 M2 server. By performing reads and writes simultaneously, it eliminates memory read-to-write blocking latency. In addition, it also offers innovative data reliability and security features to help improve data integrity, including enhanced CRC protection, data retry on error detect and buffer registers for improved fault isolation.

An important guideline is to populate equivalent ranks per channel. For instance, **mixing one single-rank DIMM and one dual-rank DIMM in a channel should be avoided**.

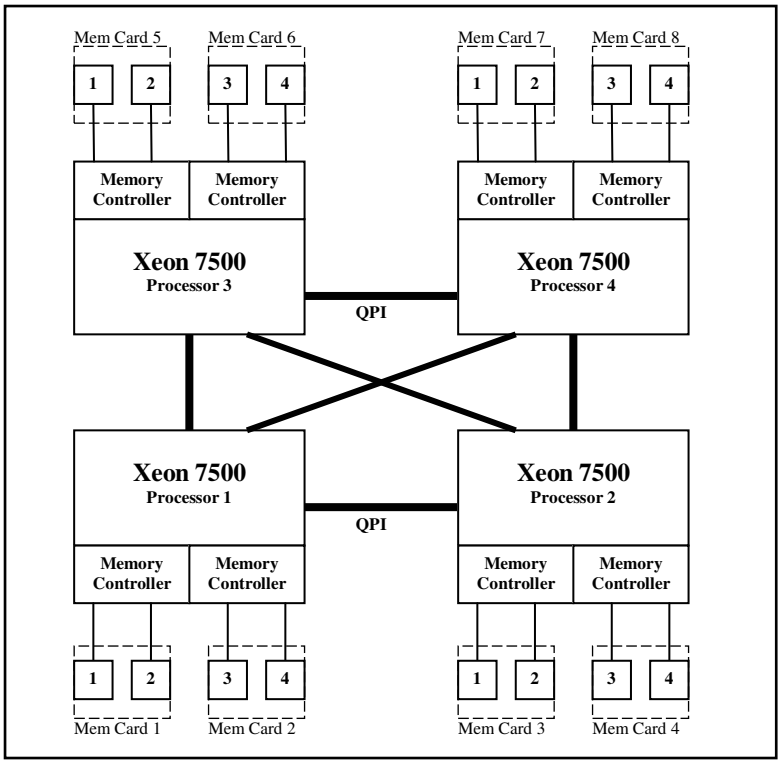

Notes: DIMMs must be installed in matching pairs. Also, each CPU requires at least 2 DIMMs. It is important to ensure that all memory channels in each processor are populated. The relative memory bandwidth decreases as the number of channels populated decreases. This is because

the bandwidth of all the memory channels is utilized to support the capability of the processor.

So, as the channels are decreased, the burden to support the requisite bandwidth is increased on the remaining channels, causing them to become a bottleneck.

For peak performance:

- Always populate processors with equal amounts of memory to enable a balanced NUMA system.
- Always populate both memory controllers with equal memory capacity.
- Ensure an even number of ranks are populated per channel.
- Use dual-rank DIMMs whenever appropriate.

Power guidelines:

- Fewer larger DIMMs (for example 8 x 4GB DIMMs vs. 16 x 2GB DIMMs will generally have lower power requirements
- x8 DIMMs (x8 data width of rank) will generally draw less power than equivalently sized x4 DIMMs

Reliability guidelines:

- Using fewer, larger DIMMs (for example 8 x 4 GB DIMMs vs. 16 x 2GB DIMMs is generally more reliable
- Xeon 7500 Series memory controllers support IBM Chipkill memory protection technology with x4 DIMMs (x4 data width of rank) and x8 DIMMs

For increased availability beyond Chipkill error correction, the eX5 servers offer additional levels of IBM Active Memory protection:

Memory ProteXion technology provides multi*chip* error protection and works in conjunction with **Chipkill** technology—which provides multi*bit* protection per chip—and standard ECC protection, to provide multi-level memory correction.

With **memory rank sparing**, two unused ranks per memory card are configured as spares. These spare ranks are used in the event of a memory failure on the same memory card.

Memory mirroring works much like disk mirroring. The total memory is divided into two channels. Data is *written concurrently to both channels*. If a DIMM fails in one of the DIMMs in the primary channel, it is instantly disabled and the mirrored (backup) memory in the other channel becomes active (primary) until the failing DIMM is replaced. With mirroring enabled, one-half of total memory is available for use. (***Note:*** Due to the double writes to memory, performance is affected.) Mirroring is handled at the hardware level; no operating system support is required.

Memory scrubbing is an automatic daily test of all system memory. It detects and reports memory errors that might be developing before they cause a server outage. Memory scrubbing and Memory ProteXion work together. When a bit error is detected, memory scrubbing determines whether the error is recoverable. If the error is recoverable Memory ProteXion technology will write the data to new location; if it is not recoverable, scrubbing sends an alert to light path diagnostics, which then notifies IBM Systems Director.

Memory is available in **2GB**, **4GB**, **8GB**, or 1**6GB** DIMMs.

High-Performance Scalability

Scalability can be achieved in two different ways. You can choose to scale “natively” by simply cabling two chassis together via the Quick Path Interconnect (QPI) ports, for **8 sockets** and **128 DIMMs**. Or, you can scale via **MAX5**, adding memory capacity for increased performance. This configuration cables a **4-socket** chassis to a MAX5 containing **32 DIMM slots**, for a total of **4 sockets** with up to **96 DIMMs**.

Up to four QPI cables connect the MAX5 to the host x3850/x3950 X5 at speeds up to **6.4Gbps** (processor-specific), providing exceptional interchassis throughput (**100GBps** aggregate).

Because of this design, an x3850/x3950 X5 server has the incredible flexibility of starting out as a **4**-socket server with up to **1TB** of memory, **seven** PCIe adapter slots, and **eight** HDD bays or **16** SSD bays, and then ***doubling*** the processors, cache, slots, and HDD bays—or by increasing only the memory capacity by **50%**—while maintaining best-in-class performance.

The MAX5 scalability and memory expansion feature is a unique IBM enhancement, not offered by other x86 server architectures.

Please see the Legal Information section for important notices and information.

Flexible Internal Storage Capacity

The eX5 servers offer a choice of internal storage, supporting up to **eight** (**2.5**-inch) **hot-swap** Serial-Attach SCSI (**SAS**) or Serial ATA (**SATA**) hard disk drives, or up to **sixteen 1.8-inch** solid-state drives (**SSDs**) per chassis (using the IBM eXFlash option).

2.5-inch SAS

- **7,200** RPMs — **500GB** (**4.0TB** maximum *per chassis;* **8TB** in a 2-chassis configuration)
- **10,000** RPMs — 73.4, 146.8, or **300GB** (**2.4TB** *per chassis;* **4.8TB** per 2-chassis)
- **15,000** RPMs — 73.4 or **146.8GB** (**1.174TB** per chassis; **2.35TB** per 2 chassis)

2.5-inch SATA

- **7,200** RPMs — 160 or **500GB** (**4.0TB**)

1.8-inch SSD

- **50GB** (**800GB** *per chassis*)
- **200GB** (**3.2TB** *per chassis – 4Q/10)*
  - **High I/O Performance**
    - Offers up to 8X more IOPS than HDDs (67/33% read/write OLTP transaction base mix);
    - Optimized for heavy mix of read and write operations, such as transaction processing, media streaming, surveillance, file copy, logging, backup / recovery, and business Intelligence
  - **Lower-Cost IOPS Performance**
    - Yields better $/IOPS: lower capacity (GB) required to achieve higher IOPS
    - Consumes less energy and produces less heat than a conventional disk drive
  - **Superior Uptime**
    - 3X the reliability of mechanical disk drives[1]
    - No moving parts to fail
    - Enterprise wear-leveling to extend life even further
  - **Full OS Support**
    - Supports all ServerProven® OSes

High IOPS SSD PCIe Adapters

- IBM **160GB** High IOPS **SS** Class SSD PCIe Adapter (**x4 Gen 1**—100K IOPS, 750MBps)
- IBM **320GB** High IOPS **SS** Class SSD PCIe Adapter (**x4 Gen 1**—100K IOPS, 750MBps)
- IBM **320GB** High IOPS **MS** Class SSD PCIe Adapter (**x4 Gen** 1—100K IOPS, 750MBps)
- IBM **320GB** High IOPS **SD** Class SSD PCIe Adapter (**x8 Gen 2**—**200K** IOPS, **1.5GBps**)
- IBM **640GB** High IOPS **MLC** Duo Adapter (**x8 Gen** 2—**200K** IOPS, **1.5GBps**)

For read-intensive transactional workloads requiring extreme IOPS performance, such as database, video-on-demand, and caching, IBM offers **High IOPS SSD PCIe Adapters**, equivalent to the IOPS output of approximately **500**[2] 3.5-inch 600GB 15K SAS HDDs, with **99% lower latency** (30 µs), and **7.7x** the bandwidth of a HDD. On a performance-per-watt basis, these adapters outperform HDDs by up to **445x**[3].

Because these adapters go in PCIe slots, rather than drive bays, they can be used *in addition* to standard HDDs and SSDs. This makes them ideal for virtualized servers running multiple tasks, requiring both write- and read-intensive storage media. As flash media they also offer a predictable lifetime, compared to the sporadic failure rate of HDDs, using N+1 chip-level redundancy and 11-bit ECC protection.

If you need more storage space, terabyte capacities are possible with external direct-attach NAS and SAN solutions.

[1] Approximately 3,000,000 hours MTBF for SSDs vs. approximately 1,000,000 hours for enterprise 3.5-inch 15K SAS HDDs.

[2] According to internal IBM testing, the typical **maximum IOPS** of a 600GB 15K **3.5**-inch hot-swap SAS HDD is ~**400** (4K random reads), while the typical IOPS of a 640GB High IOPS SSD PCIe adapter is ~**200,000**, or ~**500x**. The typical **bandwidth** of the same HDD = ~**195MBps** (64K sequential reads), while the typical bandwidth of a 640GB High IOPS MLC adapter is ~**1.5GBps**, or ~**7.7x**.

[3] 97,014 IOPS / 9W = **10,779 IOPS** per watt (160GB/320GB adapters). 196,000 IOPS / 12W = **16,333 IOPS** per watt (640GB adapters). 400 IOPS / 16.5W = **24.2 IOPS** per watt (600GB 15K 3.5-inch hot-swap SAS HDD).

Please see the Legal Information section for important notices and information.

Storage Controllers

All models include an integrated **ServeRAID-BR10i V2 SAS/SATA Controller** in a dedicated x8 PCIe slot. Supported controllers include:

- The 3Gbps[4] (x8 PCIe) **ServeRAID-BR10i V2** controller — hardware **RAID-0**/**1**/**1E** support (no cache) for up to **8** HDDs or SSDs.
- The **6Gbps** (x8 PCIe) **ServeRAID-M1015 SAS/SATA** controller — **RAID-0**/**1**/**10** (no cache) for up to **8** drives. The **IBM ServeRAID M1000 Series Advance Feature Key** adds **RAID-5** and Self-Encrypting Disk (**SED**) support.
- The **6Gbps** (x8 PCIe) **ServeRAID-M5015 SAS/SATA** controller — enhanced performance with **512MB** of cache memory and battery backup, and supports **RAID-0**/**1**/**10**/**5**/**50** for up to **8** drives. The **IBM ServeRAID M5000 Series Advance Feature Key** adds **RAID-6**/**60** and **SED** support to the M5015.
- The **6Gbps** (x8 PCIe) **ServeRAID-M5025 SAS/SATA** controller — enhanced performance with **512MB** of cache memory and battery backup. It supports **RAID-0**/**1**/**10**/**5**/**50** for up to **8** internal drives, plus multiple external IBM System Storage® expansion units (with up to **240** HDDs)**.**

Additional external SAS/SATA disk storage, as well as tape backup, is available via one of several supported iSCSI or FC SAN controllers.

Versatile Drive Bays

The x3850 X5 and x3950 X5 servers include either 1 or 2 drive backplanes (model-specific).The single-backplane models are upgradeable to 2 backplanes. Each backplane supports either **4 2.5-inch hot-swap HDDs** or, within the same space, **8 1.8-inch hot-swap SSDs**. Thus, you can achieve up to **8** HDDs, or **16** SSDs using up to *two* IBM **eXFlash** storage options. Each individual SSD is capable of sustaining approximately **4,300 IOPS**, which is more than 10X what an enterprise 15,000 RPM 2.5-inch 146.8GB SAS HDD can achieve (420 IOPS). Thus, *each* eXFlash module (containing **8** SSDs) is capable of producing up to **34,400 IOPS**, vs. only **3,360 IOPs** for **8** SAS HDDs. This means one x3850/x3950 X5, with two eXFlash options (16 drives), is capable of generating **68,800 IOPs**, vs. only about 6,700 IOPS for storage servers using an equal number of HDDs. In a 2-chassis configuration, up to **16 HDDs** or **32 SSDs** are supported (for up to **137,600 IOPS**).

Each customer-installable eXFlash option includes a drive cage, SSD backplane, and two x4 SAS/SATA cables, which can be installed in seconds as a single unit. **One** or **two** ServeRAID controllers are required *per eXFlash option*. For RAID support, use one **ServeRAID-M1015, M5014, M5015** or **M5025** controller. For enhanced IOPS throughput, use *one* **ServeRAID B5015** *non-RAID* controller, which is optimized for SSD IOPS performance. (***Note:*** Because SSDs offer triple[5] the reliability of SAS HDDs, RAID support is less necessary.) Each adapter requires a PCIe slot. For **maximum IOPS performance**, use *two* **ServeRAID-B5015** controllers per eXFlash option.

Supported *single-chassis* configurations include:

- **4** 2.5-inch **HDDs** + **8** 1.8-inch **SSDs** (with 1 eXFlash option)
- **8** 2.5-inch **HDDs**
- **16** 1.8-inch **SSDs** (with 2 eXFlash options)

Supported *two-chassis* configurations include:

- **8** 2.5-inch **HDDs** + **16** 1.8-inch **SSDs** (with 2 eXFlash options)
- **16** 2.5-inch **HDDs**
- **32** 1.8-inch **SSDs** (with 4 eXFlash options)

An optional (ultraslim, 0.5") optical drive with a SATA interface is supported in all x3850 X5 and x3950 X5 servers. An external USB floppy drive or optical drive may be used, if needed.

For still more storage, NAS, iSCSI SAN, or FC SAN external expansion options can be added, using an optional controller.

High-Performance Adapter Slots

There are **seven PCIe (PCI Express) Gen 2** adapter slots standard (per chassis). **PCI Express** Gen 2 is the latest high-performance, low-latency, next-generation serial I/O bus. **Slot 7** contains a **10Gb Ethernet Adapter** in all models except for the 7145-ARx. **Slots 1** through **4** are **full**-length/full-height. **Slots 5-7** are **half**-length/full-height. **Slot 1** is a **x16** ("by 16") slot, both physically and electrically (meaning it operates at the full **x16 16GBps**[6] bidirectional speed). **Slot 2** is physically **x8** but electrically **x4** (**4GBps**). **Slots 3-7** are **x8** (**8GBps**) both physically and electrically.

Because the **SAS**, **ServeRAID**, dual **Gigabit Ethernet**, **video**, and **systems management**

Please see the Legal Information section for important notices and information.

controllers are all integrated onto the system board, six of the seven adapter slots are *available*, which offers you a wide degree of latitude in expansion options.

10 Gigabit Ethernet Virtual Fabric Adapter for IBM

Most x3850 X5 and x3950 X5 systems include one **dual-port 10Gb Virtual Fabric** adapter. The Emulex Virtual Fabric Adapter is an industry-leading performance and scalability per watt, dual-port network adapter for 10Gbps Ethernet (10GbE) networks. It offers the benefits and flexibility of I/O convergence in a single end-to-end solution. Protocol offload for stateless TCP/IP and TCP Chimney provide maximum bandwidth with minimum use of CPU resources. It achieves line rate 10Gbps performance with support for TCP/IP stateless offloads and TCP Offload Engine (TOE) support. TOE reduces system processor utilization, providing increased system performance and reducing overall system power requirements.

The adapter is based on the Emulex OneConnect Universal Converged Network Adapter (UCNA) platform that also includes the capability for future upgrades to Fibre Channel over Ethernet (FCoE) and iSCSI protocol offloads. By using a common infrastructure for Ethernet and storage networks, data centers can reduce capital expense (CapEx) for adapters, switches and cables, and operational expense (OpEx) for power, cooling and IT administration.

End-to-end data protection with hardware parity, CRC, ECC and other advanced error checking and correcting ensure that data is safe from corruption.

Dual 10Gbps Ethernet ports:

- IPv4/IPv6 TCP, UDP checksum offload; Large Send Offload (LSO); Large Receive Offload; Receive Side Scaling (RSS); IPV4 TCP Chimney Offload
- VLAN insertion and extraction
- Jumbo frames up to 9000 Bytes
- Preboot eXecutive Environment (PXE) 2.0 network boot support
- Interrupt coalescing
- Load balancing and failover support including adapter fault tolerance (AFT), switch fault tolerance (SFT), adaptive load balancing (ALB), teaming support and IEEE 802.3ad.

Note: You must have either one SFP+ transceiver or one SFP+ direct-attached cable for *each* of the two 10Gb ports on the adapter.

Dual Gigabit Ethernet Controller

The eX5 systems include one integrated **dual-port Broadcom 5709C** Gigabit Ethernet controller, with **TOE** (TCP Offload Engine) acceleration support, as well as **load-balancing** and **failover** capabilities between the two ports.

TOE support helps improve overall system performance by offloading TCP/IP protocol processing from the system microprocessor to the onboard Ethernet TOE processor.

The controller also supports highly secure remote power management using **IPMI 2.0**, plus **Wake on LAN**® and **PXE** (Preboot Execution Environment) flash interface. Optional PCI adapters offering failover and load balancing between adapters are available for added throughput and increased system availability.

Hot-Swap/Redundant Components

System availability is maximized through the extensive use of hot-swap and redundant components, including:

- **Hot-swap redundant hard disk drives** and **solid-state drives**
- **Hot-swap redundant cooling fans** (three redundant pairs)
- **Hot-swap redundant power supplies**[7]

Ultra-Efficient Cooling

Six strategically located **hot-swap**/**redundant** fans, combined with efficient airflow paths, provide highly effective system cooling for the x3850 X5 and x3950 X5 systems, known as **Calibrated Vectored Cooling**. The fans are arranged to cool **three separate zones**, left, center, and right, with one 120mm fan each in the left and center zones and a pair of redundant 30mm fans in the third zone. The left and center zones are redundant with fans contained in the removable power supply units.

The fans automatically adjust speeds in response to changing thermal requirements, depending on the zone, redundancy, and internal temperatures. When the temperature inside the server increases, the fans speed up to maintain the proper ambient temperature. When the temperature returns to a normal operating level, the fans return to their default speed. Why not

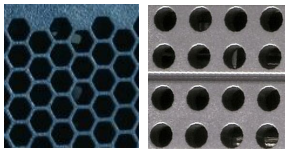

simply run the fans at 100% capacity all the time? For several good reasons: to reduce the ambient noise, reduce the wear-and-tear on the fans and reduce the server power draw. The reduction in ambient noise and power draw may be relatively minor for a single server, but if you have a larger number of systems the aggregate savings in power and noise pressure can be significant.

In addition, the server uses **hexagonal ventilation holes** in the chassis. Hexagonal holes can be grouped more densely than round holes, providing greater airflow through the system cover.

This cooling scheme is important because newer, more powerful processors generate a significant amount of heat, and heat must be controlled for the system to function properly.

Light Path Diagnostics

Light path diagnostics enables a technician to quickly identify and locate a failed or failing system component, such as a specific memory DIMM. (Monitored components include processors, memory, power supplies, fans, HDDs, adapters, and voltage regulator modules (VRMs), plus system temperature.) This enables quick replacement of the component, which helps increase server uptime and lower operating costs.

The front of the server has an LED indicator light to show possible component failures. If the front LED indicates an error condition, by pressing a button on the front of the server an LED panel will pop out and drop down for easy viewing *without the need to open the server cover* or remove the server from the rack. The light path diagnostics panel tells the servicer which component requires attention. In addition, many components have their own identifying LEDs. For example, each of the memory modules has an LED next to the socket, as do all processors, all adapter slots, all fans, all power supplies, all voltage regulator modules and the service processor, allowing the servicer to easily identify exactly which component needs servicing. By following the "light path," the component can be replaced quickly, and without guesswork. (***Note:*** In the event of a failed DIMM, the system will restart and mark the DIMM as bad while offline, thus allowing the system to continue running, with reduced memory capacity, until serviced.)

In addition to the outstanding diagnostic detection, there is a new feature with a digital readout that indicates at what stage of the boot process your system may have experienced trouble. This can reduce the troubleshooting time and help get your system up and running quickly and efficiently.

Energy-Smart Features

IBM servers work hard to keep your energy bills low—from high-efficiency power supplies and fans to lower-draw processors, memory, and solid-state drives. Technologies such as these, along with the Xeon 7500/6500 series processors that intelligently adjust their voltage and frequency, help take costs out of IT:

- **95W 8-core processors** use **27%** less energy than 130W processors.
- **1.5V DDR3 DIMMs** consume **10-15%** less energy than 1.8V DDR2 DIMMs used in older servers.
- **Solid-State Drives** consume up to **80%** less energy than 2.5-inch HDDs and up to **88%** less than 3.5-inch HDDs.
- On a performance-per-watt basis, **High IOPS SSD Adapters** outperform enterprise HDDs by up to 445-to-1[8]
- **Dynamic fan speeds** — In the event of a fan failure, the other fans will run faster to compensate until the failing fan is replaced. Competitive fans must run faster at all times, just in case, wasting power.

Other Features

- **Trusted Platform Module** — This module provides a highly secure start-up process from power-on through hand-off to the operating system boot loader. ACPI support is provided to allow ACPI-enabled operating systems to access the security features of this module. (TCG V1.2-compliant.)
- **Eight USB 2.0 ports** — Provides flexibility to add high-speed external devices. The USB 2.0 specification supports up to 480Mbps transfer rates. (Note: Not all USB 2.0 devices are capable of achieving this rate.) **Two** ports are provided on the front of the server, **four** are on the back, and **two** are internal to support a **USB 2.0 Flash Memory Key** with embedded hypervisor
- **Toolless slides** — Allows quick rack installation and quicker upgrade and servicing of the server.
- **Toolless chassis** — The cover can be opened without tools, and many components can be added or replaced without tools, including the processors[9], DIMMs, hot-swap HDDs, PCIe

adapters, an optional optical drive, and the integrated or optional ServeRAID controller. This can save a servicer significant time.

Advanced Systems Management Capabilities

The x3850 X5 and x3950 X5 have a high level of systems management capabilities that are well-suited to remote locations as well as to stand-alone environments. Features include UEFI, IMM with Virtual Media Key, IBM ToolsCenter, IBM Systems Director Active Energy Manager, Automatic Server Restart, Automatic Node Failover, Wake on LAN® support, PXE support, text console redirect, Predictive Failure Analysis, and IBM Systems Director.

The **IMM** provides industry-standard **Intelligent Platform Management Interface** (**IPMI**) **2.0**-compliant systems management. It provides a number of important system functions, including:

- Monitoring of system and battery voltage, system temperature, fans, power supplies, processor and DIMM status
- Fan speed control
- Product ID and Family ID detection
- Highly secure remote power on/off
- System reset control
- NMI/SMI detection and generation
- System diagnostic LED control (power, HDD, activity, alerts, heartbeat)
- IPMI over LAN
- Serial Over LAN
- Proxy server support
- LAN messaging and alerting
- Text console redirection over LAN
- Predictive Failure Analysis for system fans
- Web-based out-of-band control
- SSL (Secure Socket Layer) and LDAP (Lightweight Directory Access Protocol) support
- VLAN support
- Enhanced authentication and encryption algorithms (RMCP+, AES)
- Local update of IMM firmware
- Firmware firewall
- Support for IPMI v2.0 compliant management software (e.g., xCAT)
- Other mandatory and optional IPMI IMM functions

The IMM alerts IBM Systems Director to anomalous environmental factors, such as voltage and thermal conditions—even if the server has failed.

The x3850 X5 and x3950 X5 also include standard an **IBM Virtual Media Key** for additional systems management capabilities, including:

- Latest OS failure screen capture
- Graphical console redirection over LAN
- Remote virtual floppy and CD-ROM
- High-speed remote redirection of PCI video, keyboard and mouse

IBM **ToolsCenter** consolidates 42 needed tools for managing servers individually into an integrated suite of 8 tools. They are organized by function: deployment, updates, configuration and diagnostics. Tools are now simpler to access and use with a single webpage for access, a common look and feel and a common command line interface for the scripting tools. The ToolsCenter **Bootable Media Creator** offers significantly more functionality than past tools with the ability to add more tools to the bootable image and to automatically download the bootable environment if needed. Bootable Media Creator allows you to create bootable CDs, DVD, and USB keys for updates customized to your systems.

Automatic Server Restart (ASR) helps reduce downtime by restarting the server automatically in the event of a system lockup. If the operating system crashes or the hardware freezes, it automatically triggers the ASR hardware, which immediately restarts the server (and logs an ASR event with IBM Systems Director). These features are designed so that *no more than five minutes can pass before the server is restarted*.

Automatic Node Failover helps improve the availability of multi-node systems. If one node of a multi-node system experiences a failure, the remaining node will reboot and continue running.

Wake on LAN permits the server to be remotely powered on if it has been shut off. Once powered up, the server can be controlled across the network, using the **Preboot Execution Environment** (PXE).

Like Wake on LAN, PXE is system firmware. It enables software such as the optional **IBM Remote Deployment Manager** to take control of a system before the BIOS, operating system or applications are loaded (using Wake on LAN/PXE) and lets an administrator perform many low-level tasks remotely that would otherwise require a visit to each system. These tasks may include such things as formatting a hard disk drive, updating system firmware, or deploying a Windows or Linux operating system.

Text and Graphical Console Redirection support allows the administrator to remotely view system text messages and graphics over serial or LAN.

Predictive Failure Analysis (PFA) is designed to allow the eX5 servers to detect impending failure of supported components (processors, memory, PCIe slots, VRMs, power supplies, fans, and hard disk drives) *before* actual failure, and alert the administrator through IBM Systems Director. This gives you the ability to replace the failing component *before* it fails, resulting in increased uptime.

IBM Systems Director software for advanced workgroup management is included standard. IBM Systems Director comes with a portfolio of tools, including **IBM Systems Director Active Energy Manager**, **Service and Support Manager**, and others. *System Availability* (a no-charge download) and *Capacity Manager* (sold separately) are available as add-ons for additional server management and increased availability. IBM Systems Director provides a single uniform graphical interface for all of these systems management functions. IBM Systems Director enables you to customize thresholds and monitor system components (for things like temperature, voltage regulation, etc.) to help maximize uptime.

IBM developed IBM **Systems Director Active Energy Manager** to put control of system power-saving features at the fingertips of administrators. Active Energy Manager is designed to take advantage of new features, such as monitoring power usage and balancing the performance of the system according to available power input. It provides the ability to plan and predict power consumption based on your hardware configuration. It also helps enable you to reduce the infrastructure required for redundancy, by using fewer servers on smaller power feeds and potentially lowering your overall data center support costs. It does this by inventorying all components, then adding up the total power draw and tracking the usage. It also includes power management features to help administrators manage or reduce power usage.

Extensive System Support Features

The IBM services and technical support portfolio provides world-class, consistent, high-quality service and support. The eX5 servers offer a number of tools and services designed to make ownership a positive experience. From the start, IBM programs make it easier for you to plan for, configure and purchase System x or xSeries servers, get them running and keep them running long-term. These features include IBM ServerProven, the IBM Standalone Solutions Configuration Tool, IBM System x and BladeCenter Power Configurator, IBM ServerGuide, IBM Systems Director Service and Support Manager, Product Customization Services and extensive technical support offerings.

The IBM **ServerProven** program provides the confidence that specific options and operating systems have been tested on the server and are officially supported to work together. It is updated frequently to ensure that the latest compatibility information is always at your fingertips.

The IBM **Standalone Solutions Configuration Tool** (SSCT) is a downloadable tool that simplifies the often complex chore of configuring a full rack of servers (including blade servers) and confirming that you have all the cables, power distribution units, KVM (keyboard, video and mouse) switch boxes and other components you need, as well as the proper airflow clearances, electrical circuits and other environmental conditions.

IBM **System x and BladeCenter Power Configurator** helps IT managers plan for data center power needs, by providing the following information for specific configurations of System x and BladeCenter systems: *power input* (watts), *PDU sizing* (amps), *heat output* (BTUs), *airflow requirements through chassis* (CFM), *VA rating*, *leakage current* (mA), and *peak inrush current* (amps).

IBM **ServerGuide** (installed from CD) simplifies the process of installing and configuring System x and xSeries servers. ServerGuide goes beyond mere hardware configuration by assisting with the automated installation of the Microsoft® Windows® Server 2008 operating systems, device drivers and other system components, with minimal user intervention. (Drivers are also included for support of Novell NetWare, Red Hat Linux and SUSE LINUX.) This focus on deployment helps you reduce both your total cost of ownership and the complexity that administrators and technical personnel face.

IBM Systems Director Service and Support Manager (previously called IBM Electronic Service Agent™) is an innovative "call home" feature that allows System x and BladeCenter servers to automatically report hardware problems to IBM support, which can even dispatch onsite service if necessary to those customers entitled to onsite support under the terms of their warranty or an IBM Maintenance Agreement. Service and Support Manager resides on a server and provides electronic support and problem management capabilities through a highly secure electronic dialogue between your systems and IBM. It monitors networked servers for hardware errors and it can perform hardware and software inventories and report inventory changes to IBM. All information sent to IBM is stored in a highly secure database and used for improved problem determination.

Additional services include hardware warranty upgrades and factory-installed **Product Customization Services** (PCS), such as asset tagging, hardware integration, software imaging and operating systems personalization.

IBM offers extensive **technical support** by phone and via the Web. Support options include links to forums/newsgroups, problem submission, online shopping support, service offerings, device drivers for all IBM product lines, software downloads and even upcoming technical seminar worldwide schedules and registration. Also available are remote installation, configuration and usage support for System x and xSeries hardware and software, as well as onsite custom services to provide the level of expertise you require.

IBM Maintenance and Technical Support solutions can help you get the most out of your IT investment by reducing support costs, increasing availability and simplifying management with integrated support for your multiproduct, multivendor hardware and software environment. For more information on hardware maintenance, software support, solution support and managed support, visit http://**ibm.com**/services/maintenance.

Key Options

IBM options for System x servers help you take your servers to a higher level

You can rely on System x options to supply a complete solution for your business needs. Options help you create an optimized server system to meet your data protection, storage and availability needs. Every IBM option is designed and tested for peak performance and flexibility, helping to maximize your return on investment. The combination of System x servers and options lets you keep your fingers on the pulse of your e-business.

Processors — The Intel Xeon processor provides high clock rates, a choice of 8, 6, or 4 cores, 64-bit extensions, a large L3 cache and advanced features for availability and manageability. Large cache size, combined with a fast **1066MHz** memory access, reduces memory latency and facilitates the movement of data through the processor and I/O devices. (***Note:*** System performance depends not only on the number of processors in the server but also on the power and functionality of each processor.) Adding a second, third, or fourth processor may be a cost-effective way to achieve significant performance improvements.

Memory — Memory is a significant factor in systems application performance. Adding more memory to a System x server is one of the most effective ways to increase application performance. An 8-core processor should have twice as much memory as a 4-core processor. x3850 X5 and x3950 X5 servers support **64** DIMMs natively and **96** DIMMs using a **MAX5** memory expansion unit.

Hard Disk Drives — IBM hard disk drives help customers improve the transaction and cost performance of their System x servers. The choice of hard disk drives can be a critical aspect of maximizing the I/O throughput of the system. **SAS** hard disk drives (2.5-inch) are available for the eX5 servers with capacities of **500GB** at **7,200** RPMs, up to **300GB** at **10,000** RPMs, and **146.8GB** at **15,000** RPMs. Enterprise-class 2.5-inch **SATA** hard disk drives are available with capacities of up to **500GB** at **7,200** RPMs.

Solid State Drives — IBM offers a **50.0GB** 1.8" **solid-state drive** as a higher reliability, lower-energy-use alternative to internal HDDs. They can be used as a highly available boot drive, for storing disk images, or for other uses that stress read performance. A **200GB** 1.8" drive is also supported in the x3850 X5 and x3950 X5, when available.

IBM ServeRAID Controllers — System x servers using ServeRAID technology allow companies to build a reliable foundation for business-critical computing. IBM ServeRAID technology allows an array consisting of multiple physical hard disk drives to be treated as one logical drive. ServeRAID technology also allows data to be stored redundantly, across multiple hard disk drives— enhancing both the integrity and the availability of the data. SAS and SATA ServeRAID controllers offer enhanced performance due to onboard processors and cache. Because IBM ServeRAID controllers can help significantly improve data transfer rates, this technology is extremely effective when implementing demanding, transaction-oriented applications. By employing the advanced fault tolerance of IBM ServeRAID technology, companies can effectively implement networked business systems that require large amounts of storage space for data and applications that must be available for their businesses to

Please see the Legal Information section for important notices and information.

continue operating.

Supported SAS/SATA RAID controllers include the IBM **ServeRAID-M1015** (x8 PCIe, RAID-0/1/10), IBM **ServeRAID-M5015** (x8 PCIe, **512MB** cache, battery backup, RAID-0/1/10/5/6; optional RAID-50/60 and SED) and the **IBM ServeRAID-M5025** (x8 PCIe, **512MB** cache, battery backup, RAID-0/1/10/5/6; optional RAID-50/60 and SED; supports multiple external IBM **System Storage** expansion units).

External Storage — The IBM **System Storage EXP810** and **EXP3000** series expansion units, as well as the **DS3000, DS4000, and DS8000** series storage subsystems and **N3000**, **N5000**, **N6000**, and **N7000** NAS systems, comprise a powerful and broad shared storage family with integrated management software designed to meet midrange and enterprise needs.

x3850 X5 / x3950 X5 Images

Front View

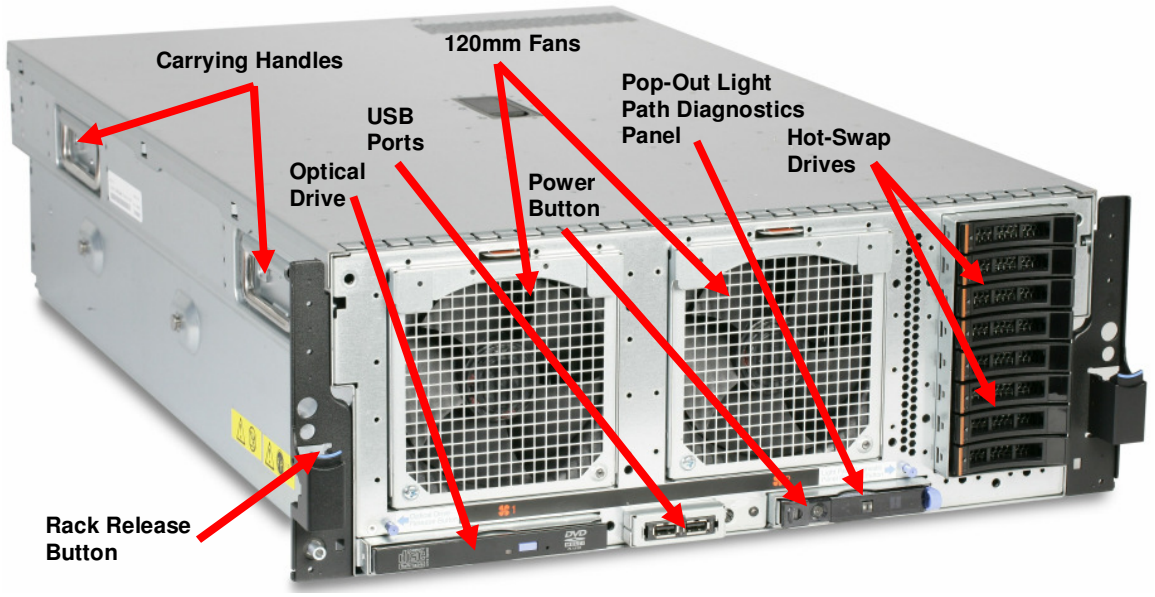

Rear View

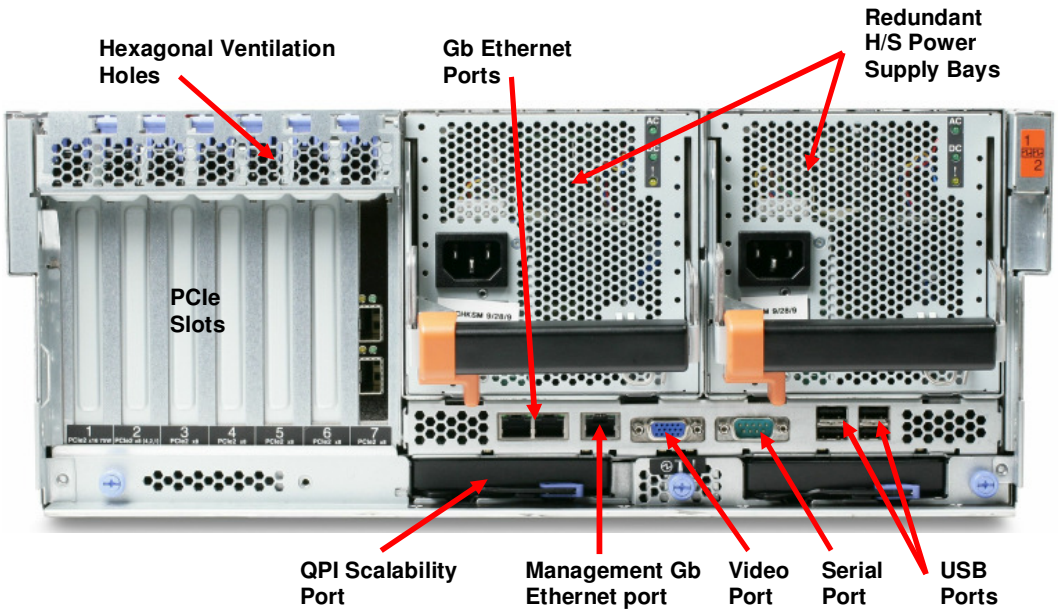

Interior View

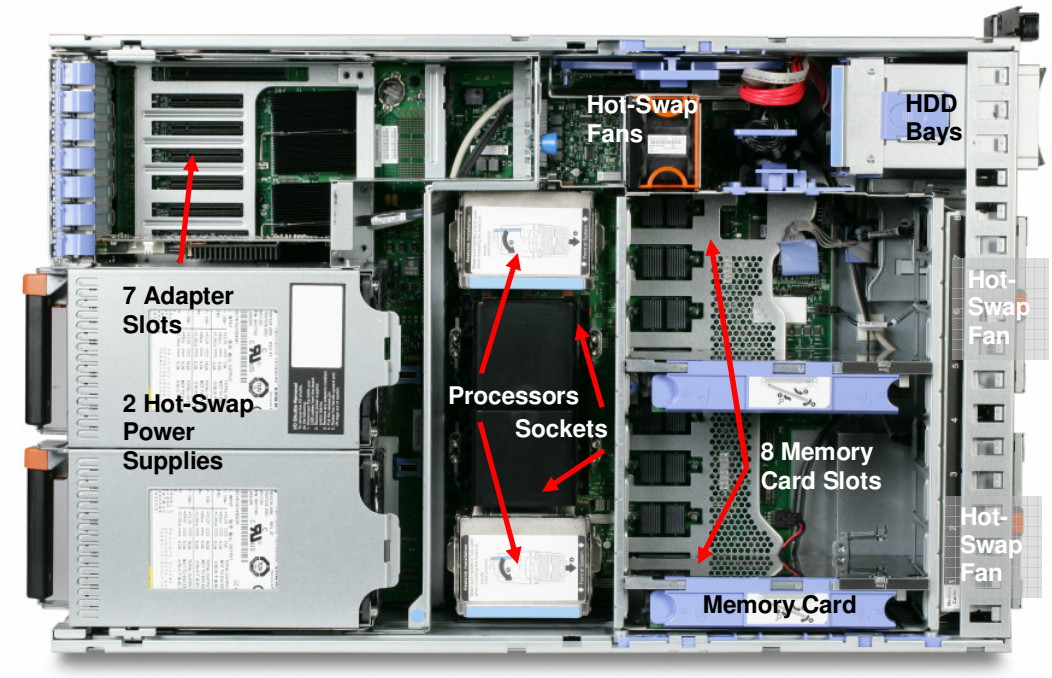

MAX5 Interior View

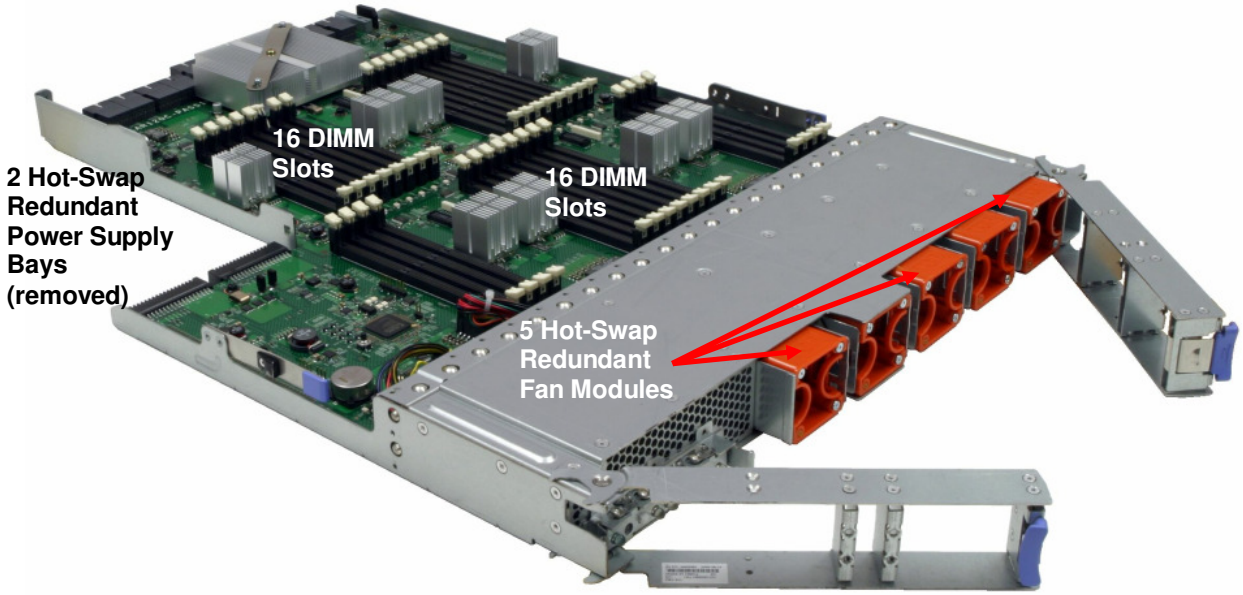

x3850 X5 / x3950 X5 Specifications			
Machine type	**x3850 X5** **7145-1Rx, 2Rx/2Sx, 3Rx, 4Rx/4Sx, 5Rx/5Sx, ARx**	**x3950 X5** **7145-4Cx/4Dx/4Fx, 5Gx**	
Form factor (standard/maximum)	4U / 10U (2 4U chassis/nodes + 2 1U MAX5 memory expansion units)		
Processor type	**Eight-core Xeon (X75xx)** **2.0GHz** X7550 (4Cx/4Dx/4Fx/4Rx/4Sx), **2.26GHz** X7560	**Six-core Xeon (E75xx)** **1.86GHz** E7530 (2Rx/2Sx), **2.0GHz** E7540 (3Rx)	**Four-core Xeon (E75xx)** **1.86GHz** E7520 (1Rx)

	(5Gx/5Rx/5Sx)			
Processor cores per chassis	**32** (4Cx/4Dx/4Fx/4Rx/4Sx, 5Gx/5Rx/5Sx)	**24** (2Rx/2Sx, 3Rx)	**16** (1Rx)	
Processor cores per 2-chassis system	**64** (4Cx/4Dx/4Fx/4Rx/4Sx, 5Gx/5Rx/5Sx)	**48** (2Rx/2Sx, 3Rx)	**32** (1Rx)	
Maximum processor power draw	**130W** (5Gx/5Rx/5Sx)	**105W** (3Rx, 4Cx/4Dx/4Fx/4Rx/4Sx)	**95W** (1Rx, 2Rx/2Sx)	
QuickPath Interconnect (QPI) speed (gigatransfers per second)	**6.4 GTps** (3Rx, 4Cx/4Dx/4Fx/4Rx/4Sx, 5Gx/5Rx/5Sx)	**5.86 GTps** (2Rx/2Sx)	**4.8 GTps** (1Rx)	
# of processors standard / maximum per chassis	**2** / **4**; up to 8 total (2 chassis)—1Rx, 2Rx, 3Rx, 4Rx, 5Rx	**4** / **4**; up to 8 total (2 chassis)—2Sx, 4Cx/4Dx/4Fx/4Sx, 5Gx/5Sx		
Internal L3 cache	**24MB** (2Rx/2Sx, 4Cx/4Dx/4Fx/4Rx/4Sx)	**18MB** (3Rx, 5Gx/5Rx/5Sx)	**12MB** (1Rx)	
Chipset	Intel 7500 chipset / IBM eX5			
Scalability ports	4 QPI Ports per chassis			
# of interconnected nodes (chassis) maximum	2			
Maximum memory per node / with 2 nodes	**1TB** / **2TB**			
Standard memory[10] (per chassis)	**384GB** (64 DIMMs x 4GB on 8 memory cards; 32 x 4Gb in MAX5)—4Cx/4Dx	**128GB** (32 x 4GB on 8 memory cards)—4Fx, 5Gx	**40GB** (8 x 4GB on 4 memory cards in chassis; 2 x 4Gb in MAX5) —2Sx, 4Sx, 5Sx	**16GB** (4 x 4GB on 2 memory cards)—1Rx, 2Rx, 3Rx, 4Rx, 5Rx
# of DIMM sockets total / available for chassis and MAX5	**96** / **0** (4Cx/4Dx)	**64** / **32** (4Fx, 5Gx)	**96** / **86** (2Sx, 4Sx, 5Sx)	**64** / **60** (1Rx, 2Rx, 3Rx, 4Rx, 5Rx)
MAX5 memory expansion unit capacity	**512GB** (32 x 16GB)			
Maximum memory for 1 node plus MAX5 expansion unit	**1.5TB** (96 x 16GB)			
Supported memory types	PC3-10600 DDR3 (**1333MHz**) ECC **RDIMMs**			
Memory interleaving	Yes (up to four-way depending on memory configuration)			
DIMM capacities supported	2GB (x8 dual-rank), 4GB (x8 quad-rank), 8GB (x8 quad-rank), 16GB (x4 quad-rank)			
Chipkill protection supported	Yes (using x4 or x8 DIMMs)			
Memory ProteXion supported	Yes			
Memory rank sparing supported	Yes			
Memory mirroring supported	Yes			
Memory scrubbing supported	Yes			
# of HDD/SSD drive bays total / available	**8** / **8 2.5-inch HDDs**; or **16** / **16 1.8-inch SSDs**			
# of 5.25" bays total / available	**1** / **1** (optional optical drive with SATA interface)			
Maximum HDD/SSD capacity	**2.5-inch SAS** **4.0TB** (8 x 500GB) **hot-swap**	**2.5-inch SATA** **4.0TB** (8 x 500GB) **hot-swap**	**1.8-inch SSD** **800GB** (16 x 50GB) **hot-swap** (16 x 200GB when available)	

Internal HDD/SSD capacities supported	**2.5-inch SAS** **500GB** — **7.2K** RPMs; 146.8, **300GB** — **10K** RPMs; 73.4, **146.8GB** — **15K** RPMs	**2.5-inch SATA** 160, **500GB** — **7.2K** RPMs	**1.8-inch SSD** **50GB** (200GB when available)
HDD / SSD capacity standard	**1.8-inch SSD** **800GB** (16 x 50GB)—4Fx, 5Gx	**2.5-inch SAS/SATA** None—all other models	
Additional Internal SSD storage via PCIe adapter	**160GB** High IOPS **SS** Class SSD PCIe Adapter **320GB** High IOPS **SD** Class SSD PCIe Adapter **320GB** High IOPS **MS** Class SSD PCIe Adapter		
# of HDDs standard	None		
# of optical drives standard	**None** (optional optical drive in dedicated 5.25" UltraBay)		
Embedded USB Flash Key available	Optional		
Disk drive technology	Hot-swap SAS/SATA		
External disk drives supported	Yes, via optional adapter		
External disk storage devices supported	System Storage DS4200, DS4700, DS4800, DS5100, DS5300, DS8100, DS8300, N3300, N3600, N3700, N5200, N5300, N5500, N5600, N6040, N6060, N6070, N7600, N7700, N7800, N7900 storage subsystems		
Integrated RAID controller / cache	**ServeRAID-BR10i V2** (x8 PCIe, RAID-0/1/1E)		
Optional RAID controllers supported	**ServeRAID-M1015** (x8 PCIe, no cache, RAID-0/1/10)	**ServeRAID-M5015** (x8 PCIe, 512MB cache, battery backup, RAID-0/1/10/5/6; optional RAID-50/60 and SED)	**ServeRAID-M5025** (x8 PCIe, 512MB cache, battery backup, RAID-0/1/10/5/6; optional RAID-50/60 and SED; support for multiple external IBM **System Storage** expansion units
# of adapter slots total / available	7 / 7 (ARx)	7 / 6 (all other models)	
# of PCIe physical x16/electrical x16 Gen 2 slots (16GBps)	**1** full-height/**full**-length per chassis/node		
# of PCIe physical x8/electrical x8 Gen 2 slots (8GBps)	**2** full-height/**full**-length / **3** full-height/**half**-length per chassis/node		
# of PCIe physical x8/electrical x4 Gen 2 slots (4GBps)	**1** full-height/**full**-length per chassis/node		
# of video ports	1		
Video controller	Matrox G200eV graphic controller integrated in IMM (optional nVIDIA 3D DCV graphics card supported)		
Video memory	16MB DDR2 SDRAM standard (shared with IMM)		
Maximum video resolution at 32-bit color	**1024x768** resolution (analog), with a color depth of 32 bits at **100Hz**; **1280x1024** resolution (digital), with a color depth of 32 bits at **75Hz**		
10 Gigabit Virtual Fabric Adapter	Dual-port Emulex NIC integrated (all models except ARx) per chassis/node		
Gigabit Ethernet controller	Dual-port Broadcom BCM5709C		
Network acceleration	TOE (1Gb and 10Gb)		
Failover / load-balancing-capable	Yes / Yes		
# of Gigabit Ethernet ports	2 (rear)		
# of serial ports	1 (rear)		
# of USB 2.0 ports	**8** (**2** front, **4** rear) ports, plus **2** *internal* USB connectors for hypervisor		

Integrated systems management controller	IMM	
Light path diagnostics support	Processors, memory, power supplies, fans, HDDs, adapters, voltage regulator modules (VRMs), and system temperature, with external pop-out/drop-down panel	
Predictive Failure Analysis support	Processors, memory, PCIe slots, VRMs, power supplies, fans, hard disk drives, and solid-state drives	
Power supply size (chassis / MAX5)	**1975W** universal, autoswitching, hot-swap (rear access)—chassis	**675W** universal, autoswitching, hot-swap (rear access)—MAX5
# of power supplies standard / maximum (chassis / MAX5)	**2** / **2** (chassis)	**1** / **2** (MAX5)
Hot-swap/redundant power supported (chassis / MAX5)	Yes / Yes	
# of fans/blowers standard / maximum	**6** / **6** (chassis)	**5** / **5** (MAX5)
Hot-swap/redundant fans supported (chassis)	Yes	
Rack mount method	Slides and Cable Management Arm (provided standard)	
Maximum altitude	7,000 ft; 2,133 m	
Operating temperature range	50–95º F; 10–35º C (up to 3,000 ft / 914.4 m)	50–90º F; 10–32º C (3,000 ft to 7,000 ft / 914.4m to 2,133m)
Heat emitted: typical / maximum (chassis / MAX5)	3753 / 6739 **BTUs**; 1100 / 2150 **Watts** (chassis)	??? **BTUs**; ??? **Watts** (MAX5)
Dimensions (chassis / MAX5)	6.8" (172.8mm) **H** 17.5" (444mm) **W** 28.04" (712mm) **D**	1.75" (44.4mm) **H** 17.5" (444mm) **W** 28.35" (720.2mm) **D**
Weight (chassis / MAX5)	77 – 110 **lbs** (maximum) 35.3 – 49.9 **kg**	28.2 – 33.8 **lb** (maximum) 12.8 – 15.1 **kg**
Operating systems supported	Microsoft Windows Server 2008 / R2 (Standard/Web/Enterprise/Datacenter Editions) 64-bit; Windows HPC Server 2008; RHEL 5 Server 64-bit (with and without Xen); SLES 10/11 64-bit (with and without Xen); VMware ESXi/vSphere Hypervisor 4.0	
Length of limited warranty	**3** years (parts and labor)[11] (**7145**)	**4** years (parts and labor) (**7146**)

The Bottom Line

The eX5 servers are extremely powerful systems, incorporating leading-edge industry-standard features and adding IBM-unique innovations:

Price/Performance

- **High-throughput processors** — **2.0** to **2.26GHz 8-core**, **1.86** to **2.0GHz 6-core**, or **1.86GHz 4-core** Xeon processors; up to **32** (8-core), **24** (6-core), or **16** (4-core) **processor cores** *per chassis*; up to **64** (8-core), **48** (6-core), or **32** (4-core) **processor cores** *per 2-chassis configuration*
- **Large L3 cache** — **24MB**, **18MB**, or **12MB** of external processor cache (model-specific)
- **Huge memory capacity —** Up to **1TB** in standalone configuration; up to **1.5TB** with **MAX5**; up to **2TB** in 2-node configuration
- **Fast memory access** — **1066MHz**, **977Mhz**, or **800 MHz** (processor-specific)
- **Fast memory** — **PC3-10600 DDR3** ECC memory standard with up to **four-way interleaving** and a **low-latency memory controller**
- **Embedded virtualization**
- **Fast disk technology** — **Serial-Attach SCSI** (SAS) controller in a dedicated slot
- **Fast SSD technology** — Optional SSD-optimized controller and drives
- **High-IOPS technology** — Optional **High IOPS** SSD adapters
- **Fast and flexible communications** — **Integrated 10Gb Virtual Fabric Adapter** and **dual-port Gigabit Ethernet** controller, supporting **load-balancing** and **failover**, as well as **TOE**,

RDMA, and **iSCSI** acceleration

- **Fast I/O** — **PCIe x16** (**16Gbps**) and **x8** (**8Gbps**) **Gen 2** adapter slots

Flexibility

- **Scale up or scale out** — Scalable on demand, from **2** to **8 processors**, **8** to **64** processor cores, **7** to **14** adapter slots, and **32** to **128** DIMMs (using *two* chassis); or **2** to **4** processors, **8** to **32** processor cores, **7** adapter slots, and **32** to **64** DIMMs (using one chassis and **MAX5**)
- **FlexNode partitioning** — Allows dynamically repartitioning the system at the node level
- **Huge memory capacity** — **1TB** of DDR3 memory, with **64** DIMM slots per chassis , plus another **512GB** in **32 DIMM** slots in a **MAX5** expansion unit; expandable to **2TB** of RAM, using **QPI scaling** and **2** chassis, or **3TB** using **2** chassis and **2 MAX5** units
- Integrated **flash** drive, preloaded with a hypervisor for "instant" virtualization (standard on selected x3950 X5 models, optional on others)
- A choice of **8 2.5-inch hot-swap SAS/SATA** HDDs or **16 1.8-inch hot-swap** SSDs internally (per chassis); up to **16 HDDs** or **32 SSDs** in a 2-chassis configuration)
- Optional IBM **160GB, 320GB**, or **640GB High IOPS SSD Adapters** for extremely high IOPS and bandwidth storage
- **High-performance external I/O expansion** — **Eight** 480Mbps **USB 2.0** ports (two front, four rear, two internal)
- Hardware-based **RAID-0/1/1E** support standard in most models; optional slotless RAID support for **RAID-10/5/50/6/60** (up to two **ServeRAID** controllers in a multi-chassis configuration)
- **Six** *available* **PCIe** slots (1 x 16GBps, 5 x 8GBps, 1 x 4Gbps) *per chassis* in most models, (**12 total slots** in a 2-chassis configuration)
- Optional optical drive

Manageability, Serviceability and Availability

- IBM **Systems Director** systems management software, including:
  - IBM Systems Director Active Energy Manager
  - IBM Service and Support Manager
- **Integrated Management Module** (IMM):
  - **IPMI 2.0** compliance, including highly secure remote power control
  - **Text console redirection** systems management standard
  - Supports **LDAP** and **SSL** industry standards
- **Active Memory protection:**
  - Advanced **Chipkill** ECC memory correction
  - **Memory ProteXion**
  - **Spare rank memory**
  - **Memory mirroring**
  - **Memory scrubbing**
- Hardware-based **RAID-1/1E** disk mirroring standard in a dedicated slot; optional **RAID-10/5/50/6/60** highly available arrays
- **Ultra-efficient cooling** incorporating **Calibrated Vectored Cooling** features
- **Hot-swap HDDs** or **SSDs**
- **Hot-swap/redundant power supplies**
- **Hot-swap/redundant cooling**
- **Light path diagnostics** (front LED panel, drop-down light path panel)
- **PFA support for** processors, memory, PCIe slots, VRMs, power supplies, fans, hard disk drives, and solid-state drives
- **Toolless chassis** and **toolless slide** design; integrated **Cable Management Arm** simplify installation and removal of the server from the rack

Server Comparison Chart

The following table shows the suggested uses for the respective IBM System x rack-optimized servers, including comparisons of the uses for which each server is best suited:

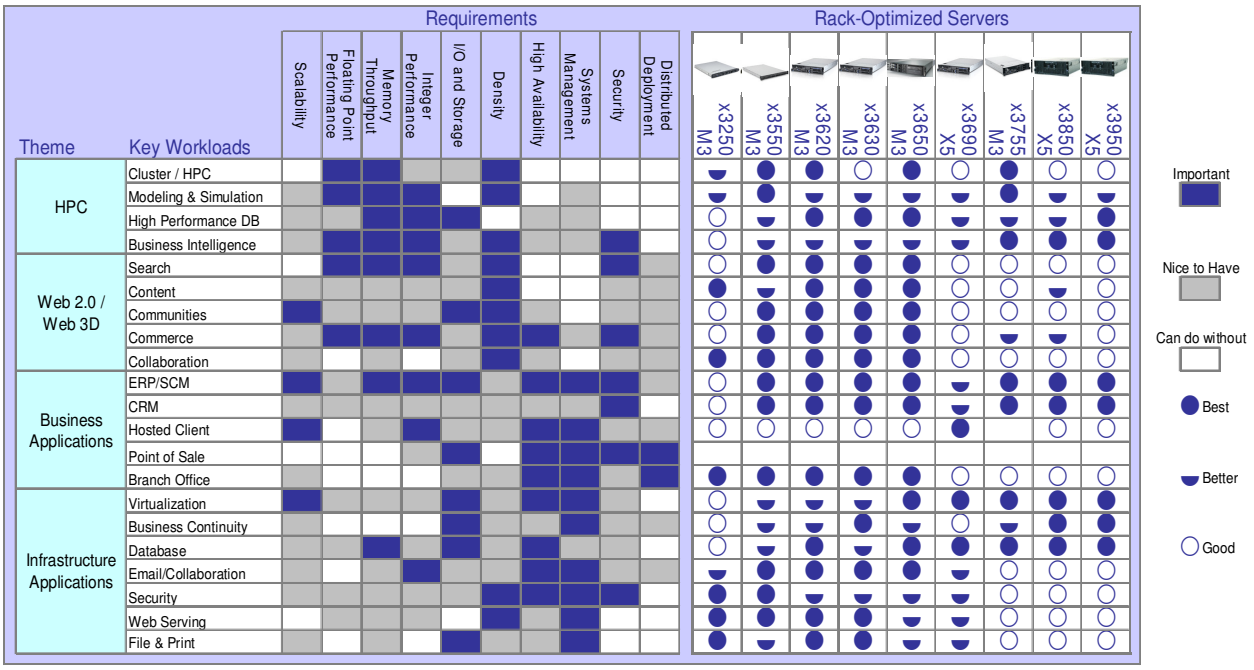
Requirements
Rack-Optimized Servers
Theme
Key Workloads
Scalability
Floating Point Performance
Memory Throughput
Integer Performance
I/O and Storage
Density
High Availability
Systems Management
Security
Distributed Deployment
x3250 M3
x3550 M3
x3620 M3
x3630 M3
x3650 M3
x3690 X5
x3755 M3
x3850 X5
x3950 X5
HPC
Cluster / HPC
Modeling & Simulation
High Performance DB
Business Intelligence
Web 2.0 / Web 3D
Search
Content
Communities
Commerce
Collaboration
Business Applications
ERP/SCM
CRM
Hosted Client
Point of Sale
Branch Office
Infrastructure Applications
Virtualization
Business Continuity
Database
Email/Collaboration
Security
Web Serving
File & Print
Important
Nice to Have
Can do without
Best
Better
Good

For More Information

IBM System x Servers	http://**ibm.com**/systems/x
IBM Systems Director Service and Support Manager	http://**ibm.com**/support/electronic
IBM System x and BladeCenter Power Configurator	http://**ibm.com**/systems/bladecenter/resources/powerconfig.html
IBM Standalone Solutions Configuration Tool	http://**ibm.com**/systems/x/hardware/configtools.html
IBM Configuration and Options Guide	http://**ibm.com**/systems/x/hardware/configtools.html
IBM ServerProven Program	http://**ibm.com**/systems/info/x86servers/serverproven/compat/us
Technical Support	http://**ibm.com**/server/support
Other Technical Support Resources	http://**ibm.com**/systems/support

Legal Information

© IBM Corporation 2010

IBM Systems and Technology Group
Dept. U2SA
3039 Cornwallis Road
Research Triangle Park, NC 27709

Produced in the USA
October 2010
All rights reserved

Warranty Information: For a copy of applicable product warranties, write to: Warranty Information, P.O. Box 12195, RTP, NC 27709, Attn: Dept. JDJA/B203. IBM makes no representation or warranty regarding third-party products or services including those designated as ServerProven or ClusterProven®. Telephone support may be subject to additional charges. For onsite labor, IBM will attempt to diagnose and resolve the problem remotely before sending a technician.

IBM, the IBM logo, ibm.com, Active Memory, Chipkill, DB2, Memory ProteXion, OnForever, ServeRAID, System Storage, Systems Director Active Energy Manager, System x, and XpandOnDemand are trademarks of IBM Corporation in the United States, other countries, or both. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available at http://ibm.com/legal/copytrade.shtml.

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Linux is a registered trademark of Linus Torvalds.

Microsoft, the Windows logo, Hyper-V, SQL Server, and Windows are trademarks or registered trademarks of Microsoft Corporation.

Other company, product and service names may be trademarks or service marks of others.

IBM reserves the right to change specifications or other product information without notice. References in this publication to IBM products or services do not imply that IBM intends to make them available in all countries in which IBM operates. IBM PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions; therefore, this statement may not apply to you.

This publication may contain links to third party sites that are not under the control of or maintained by IBM. Access to any such third party site is at the user's own risk and IBM is not responsible for the accuracy or reliability of any information, data, opinions, advice or statements made on these sites. IBM provides these links merely as a convenience and the inclusion of such links does not imply an endorsement.

Information in this presentation concerning non-IBM products was obtained from the suppliers of these products, published announcement material or other publicly available sources. IBM has not tested these products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

MB, GB and TB = 1,000,000, 1,000,000,000 and 1,000,000,000,000 bytes, respectively, when referring to storage capacity. Accessible capacity is less; up to 3GB is used in service partition. Actual storage capacity will vary based upon many factors and may be less than stated.

Performance is in Internal Throughput Rate (ITR) ratio based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput that any user will experience will depend on considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput improvements equivalent to the performance ratios stated here.

Maximum internal hard disk and memory capacities may require the replacement of any standard hard drives and/or memory and the population of all hard disk bays and memory slots with the largest currently supported drives available. When referring to variable speed CD-ROMs, CD-Rs, CD-RWs and DVDs, actual playback speed will vary and is often less than the maximum possible.

XSO03099-USEN-02